Places for Lightweight Group Meetings: The Design of Come Together

Yibo Sun and Saul Greenberg
Department of Computer Science, University of Calgary
Calgary, Alberta, Canada T2N 1N4
{saul.greenberg,yibsun}@ucalgary.ca

ABSTRACT
Lightweight group meetings are opportunistic, ad hoc, or lightly planned gatherings characterized by the informal nature of their members and their tasks. Critically, they must be very easy to set up and maintain over time. We contribute the design of Come Together, a system that supports lightweight, persistent meetings between distance-separated people. Its design is theoretically motivated by the Locales Framework, with features derived from the best of Instant Messengers and the Community Bar. Its main motivation is that any action must be simple and fast to do if it is to support lightweight group meetings. In particular, Come Together represents both people and their things as *media items*, which – unlike prior systems – can be quickly brought together to form an ad hoc *place*. Places, which are persistent, can be presented in a variety of forms (e.g., as a stand-alone window, or as an element in a sidebar), with interaction mechanisms that let a person quickly adjust the degree of awareness he or she wishes to maintain of the place and its contents. Somewhat akin to buddy lists, a *console* collects all people, artifacts, and places, where users can select them to rapidly compose meeting places.

Categories and Subject Descriptors
K.4.3 [Computers and Society]: Organizational Impacts - Computer-supported collaborative work.

General Terms
Design, Human Factors

Keywords
Ad hoc groups, awareness, casual interaction, groupware, places.

1. INTRODUCTION
People constantly come together in small groups. Typically, group members know each other to various extents, and have a real desire – driven by social, work, play or a variety of other purposes – to interact over time. While some formal gatherings happen via scheduling and planning, the vast majority are lightweight. That is, they are ad hoc, informal, casual, mostly unplanned, opportunistic, and oftentimes brief. Even so, such meetings are critical for small group effectiveness [10] [19].

Seven elements interplay to contribute to why lightweight meetings work so easily in the everyday co-located setting.

- *Informal awareness*: knowledge about presence, activity, and availability of each other triggers meeting opportunities [10].
- *Lightweight engagement.* People can quickly decide whether it is appropriate to engage into conversation by physical acts such as eye contact and body language, by greetings, and by how all move closer to one another to initiate engagements.
- *Low-cost meetings*. People can then engage with others at negligible cost, which means that even the briefest casual interaction can be done in a lightweight and routine manner.
- *Artifact sharing*, where personal artifacts (such as documents) can be brought into the meeting as a matter of course, is done simply by bringing the artifact to the attention of others and exploiting the tools ready to hand [17].
- *Proximity* of collaborators and their artifacts within a *social place* or *local* makes all the above happen within a context infused with social meaning [10] [2] [9] [4].
- *Persistence*, where people and artifacts related to a series of ongoing meetings may persist as a matter of course; i.e., prior meeting contexts can be re-established at a later time [7].
- *Ad hoc membership and involvement*, where people affiliated with the group can not only come and go, but have various degrees of involvement with it.

The challenge is how to support such lightweight meetings between distance-separated people. This is where the action has been in the design of various distributed real-time groupware systems. However, while the most successful groupware systems are excellent at supporting some of the above elements, they often fail at supporting others (see related work). Our particular interest – and the focus of this paper – is in considering how groupware can be designed to support lightweight distributed group meetings. Specifically, we concentrate on meetings and group membership that are somewhat impromptu and ad hoc, and where the level of people's involvement and likelihood of continuing a meeting in the future are highly variable.

We tackle lightweight distributed group meetings as a design challenge. First, we use related work to inform our overall design. In particular, we consider the successes and failings of Instant Messengers (e.g., [13]) and of the Community Bar [11] [15] in terms of addressing the above elements, and then we recast these elements in terms of the requirements as set out by Fitzpatrick's ethnographically-informed Locales Framework [4]. Second, we present the main features of our Come Together system, which was expressly designed to provide rapidly constructed places that ease how a group enters into lightweight meetings. As we will see, the key design points and contributions of Come Together are that: (1) both people and artifacts are treated somewhat equally, (2) both people and artifacts can exist outside of a defined group

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
GROUP'10, November 7–10, 2010, Sanibel Island, Florida, USA.
Copyright 2010 ACM 978-1-4503-0387-3/10/11...$10.00.

space, (3) places can be created simply by bringing people and/or artifacts together, (4) people and artifacts can co-exist in multiple group places, (5) the personal salience of people and items within a place can be adjusted, and (6) the buddy list idea of Instant Messengers was extended to capture not only people, but artifacts and places.

2. RELATED WORK / DESIGN PREMISES

Our requirements analysis is based on theory (the Locales Framework), and practice (successes and deficiencies of select groupware systems).

2.1 Theory: The Locales Framework

The Locales Framework [4] is a theoretical framework developed to help an analyst understand the nature of social and collaborative interactions. Here, we selectively emphasize a few of its key aspects.

Locales: Social Worlds, Sites and Means, A *social world* is a group of people with collective goal(s) and/or interest(s). Such goals do not have to be well-developed or completely knowable. Members of a social world perform their collaborative activities at a *site* (a physical and/or virtual place) furnished with *means* (resources and artifacts): collectively, this is known as a *locale.* This is somewhat akin to the concept of *place* as described in [9] and [2]. *Sites* range considerably, e.g., a physical or virtual place where people engaged in their activity, or even an abstract setting that frames the way the social world meets opportunistically as they pursue their activities. In the digital realm, sites may be explicit (e.g., a metaphor of a physical place such as a room [7]), or implicit (e.g., how a group uses a loose collection of tools to realize their interactions). The *means* are the resources, communication tools, and artifacts provided by the site and/or brought into the site. In a physical setting, resources could include tables, chairs, pen and paper, whiteboard, projector, etc. In a virtual setting, these could be the digital resources packaged as part of a tool (e.g., file transfer and VOIP capabilities in instant messengers). Artifacts include anything brought into the space and/or worked on by the group, such as documents and sketches.

Center/Periphery Principle. Most groupware treat group membership as a binary choice (i.e., member or non-member). In reality, membership of people within a social world is highly nuanced, where it is defined by a *centre-periphery continuum*. At the *centre* is the context that holds the social world together, for example, the collective goal. Core members with high levels of engagement and/ or immediate activity in the goal are typically located at this centre. Yet other members may be at the periphery, where they have a lesser level of engagement. Of course, this level of engagement can fluctuate. Thus membership at any moment of time is defined by a person's engagement on the centre/periphery continuum. In other words, actively participating people are closer to the centre of the social world, while those less engaged are further away from the centre. Similarly, artifacts can drift over the center/periphery spectrum, i.e., with artifacts relevant to the current collective goal at the centre, and secondary or stale artifacts nearer to the periphery.

Mutuality and Awareness. *Mutuality* is the mutual provision and perception of the *awareness* information exuded by people, artifacts, and interactions within the locale. In a locale, people need to be aware of the state of the social world, the site, and the means, as this helps them maintain a sense of a shared place. It also helps people move into interactions and to work with others as the interactions proceed. Indeed, such awareness – of people and / or the state of the local and the artifacts within it – has been well-defined by others as a critical element of successful group interactions ([3] [8] [10] [19] [17]).

Mutuality teases awareness into two parts: provision of the information (about the members, site and means in a locale) to others, and perception that is received by others. In turn, the centre-periphery principle affects the relationship between the provision and perception of awareness projection. A person or artifact at the centre will likely exude more information about oneself to the rest of the world, while those at the periphery will exude less. Similarly, people drawn to this centre will perceive this information at greater salience and fidelity *vs.* those at the periphery. Thus we can expect the provision and perception of awareness information to degrade selectively according to how people and things move from the center to the periphery.

Other Aspects. The Locales Framework includes several other aspects. *Individual views* acknowledge that individuals can be participating in many locales, and that they can have their own particular view of it that differs from others. *Civic structures* concern how interactions fit within a broader communal level, e.g., the relation of a locale and the things it contains to other locales. *Interaction trajectories* concern how courses of action within a local evolve over time.

Locales *vs.* Lightweight Distributed Group Meetings. Almost by definition, the elements that interplay to make lightweight group meetings work must occur as part of a very easy process on the part of its participants. This is counter to the often heavy-weight working protocols required by groupware support for formal meetings. Table 1, for example, elaborates further properties comparing example characteristics of a heavy-weight vs. a light-weight social world.

While the locales framework does not use the terms light vs. heavy weight practices, the application of locales to lightweight distributed meetings can be considered a consequence of how a social world is structured and supported. First, the creation and maintenance of a place – the locale – must be very easy and opportunistic, and it should be simple for people to become part of that local. Second, people should be able to incorporate various means, comprising both existing and new tools and artifacts, into their locale at any time and at low cost. Third, members should be able to engage within that site anywhere on the center-periphery continuum in a way that fits their current level of involvement. Fourth, mutuality – a person's sense of awareness – must reflect individual needs with little effort on the part of the person. Fifth, people and their things may be in multiple, possibly overlapping locales, and this too should be easily supported by any system. Sixth, the subtle but potentially complex aspects of the social world (as listed in Table 1) should be done at very low cost where, in our point of view, it should be an implicit consequence of the way people interact in the locale rather than as features explicitly enforced and maintained by the system. Finally, members and means can overlap and/or can exist across multiple locales. It is these overlapping relationships that define the global context of multiple locales. While complex, nuanced, and ever-shifting, this reflects what happens in the real world collaboration.

Table 1. Light vs. heavy-weight social world practices

Properties	Heavyweight	Lightweight
Structure	Formal, hierarchical	Flexibly articulated and maintained, idiosyncratic
Culture	formal	Casual and/or formative
Goal, focus, task, centre	Contrived, knowable, stable	Impromptu, unknowable, unstable
Formation and termination	Planned, created, terminated deliberately	Spontaneous, improvised, formed and dissolved spontaneously
Membership	Stable, clearly defined, hard boundary, more centered, more cohesive	Unstable, open, soft boundary, more peripheral, less cohesive
Member Roles	Managed, assigned	Self-selected, socially negotiated upon, flexible, adaptive
Access & Security	Rigid, closed	Flexible, open

2.2 Groupware for Lightweight Meetings

A subset of today's groupware already supports particular aspects of lightweight groups and their meetings. Most are specialized for particular kinds of exchanges. Email, social networking, internet forums, issue trackers, electronic calendars, project management are popular for casual communication, informal relations and some degree of artifact sharing, but don't support the real-time aspects of lightweight meetings. Instant messaging, voice over IP, video conferencing, shared screen systems, and others are excellent at supporting real time communication, but don't have notions of persistent places. Some groupware are tools (e.g., a groupware drawing tool) that could be part of the means within a locale, but do not contribute to locale creation in of itself. Indeed, few of the above systems explicitly support locales. Instead, locales are often implicit. They would be created in participant's mind as a consequence of the interplay of the groupware systems used to support group activity (but see [4],[18]).

While myriads of groupware exist, due to space limitations we focus on two systems: instance messengers, and the Community Bar. Both deserve special attention in terms of their strengths and weaknesses in supporting lightweight meetings.

Instant Messengers (IMs) work because they provide an extremely light-weight way to enter in casual interactions (e.g., see [13]). A person can quickly collect people as buddies, and see the online status of these intimate collaborators via buddy lists. If a buddy appears present, a person can opportunistically move into a textual chat with a single button click. Once engaged, both can move into richer communication (e.g., VoIP), and quickly exchange document artefacts via file transfer. While some have groupware tools to pursue actual work (e.g., shared sketchpads, screen sharing), their use is much rarer.

IM is at its best for transient conversational diads, but is limited in how it supports on-going locale formation and lightweight meetings within them. There is little or no notion of a broader group. Persistence is not supported except for chat transcripts. Tool use and artifact sharing (if any) is mostly 'out of band', and does not persist beyond the conversation. IM 'locales' are transitory. Support for center / periphery activity doesn't really exist. Text chats have hard walls, i.e., others outside the chat have no awareness of it and thus loose opportunities to join in unless explicitly invited.

Community Bar (CB) was explicitly designed around the Locales Framework [11], which is why it deserves special attention. CB presents itself via a sidebar metaphor [1]. Essentially, CB lets people create and gather into multiple electronic places (a digital realization of a local), where a place is represented as a collection of small *media items* – containers of groupware content – on the sidebar. Multiple places can be displayed within this side bar. Once a person joins a place, they are visible via a *presence media item*, i.e., others can see them in that place as a low frame rate video thumbnail. Artifacts can be posted and manipulated within a place via specialized media items. For example, a *photo item* lets people bring in a collection of photos, a *web item* lets people bring in a web page, and a *screen sharing item* lets a person show others a thumbnail of their current screen [17]. Other means affording interpersonal communication can be brought into that place, for example a *chat item* for text chats. All items in a place persist. Using a drill-down approach, people can explore any media item, where they mouse over a thumbnail to raise a larger and more detailed Tooltip Grande [1], and then engage into more direct interaction by double clicking it. This recreates that item in a full sized, feature-rich window. Collectively, the media items in a place are updated after any change, so that people stay aware of locale's activities simply by glancing over at the thumbnails.

A field evaluation of Community Bar [15] revealed that many of its design features were successful in how they realize the Locales Framework. Yet several critical deficiencies also appeared. First, the way people created places proved too heavy-weight: CB demanding *a priori* place creation along with a multi-step invitation protocol. What actually happened was that the broader community often created a single large explicit place (an 'uber local'), which was populated with a mix of many (implicit) locales as defined by the way people viewed a subset of media items within it. The consequence was that most people were interested in only a subset of people and activities in the uber-locale, with things outside their interest being annoying or distracting. This also led to a lesser feeling of belonging by those somewhat in the periphery. Second, media items containing artifacts and communications could only be created within the context of a single place. That is, a place had to be defined before an artifact media item could be created and used, and that media item could only exist in one place. Yet in real-life practice, people have individual artifacts that exist outside of places, and just bring them into one or more places as opportunity and needs warranted.

2.3 Design Considerations

The general problem we are interested in is that our current technologies do not support the actual working practices of light-weight distributed groups whose members are not subject to strict restrictions, rigid controls, and formal procedures. (e.g., [15]). To address this problem, we created a new system called Come Together (CT). CT requirements are grounded in the Locales framework. Its particular design is loosely modeled after the best of Instant Messengers and Community Bar, but adds the following design criteria to address their limitations.

1. People should be able to easily create locales (places) and allow others to spontaneously join and leave the locale in a

very light-weight manner. To achieve this, a CT Console presents, in a manner similar to a buddy list, all people and artifacts as media items, and places. Place creation is simply done by dragging selected media items of people and artifacts from the console into close proximity on the desktop: a place is then automatically formed around them. Of special note is that media items of both artifacts and people are treated equally in terms of how they are presented to a user, and in the interaction techniques people use to manipulate them.

2. People should be able to easily incorporate existing or new artifacts into places. To achieve this, we let people create an artifact as a media item both outside and inside of places, where they can bring that media item into other places simply by dropping it in there.
3. People should be able to adjust their involvement in a place (i.e., their center/periphery relationship) which in turn changes the awareness they receive. To achieve this, we adapt the sidebar approach of the Community Bar to create a multi-tiered method of viewing media items within a place. We also add different configurations of places, each tailored to different levels of participant involvement.

As we will see, these new design elements contribute significant enhancements over prior systems in terms of how they support ad hoc group formation and work practices over time.

3. An Overview of Come Together

This section introduces the main features and building blocks of Come Together. It does this by deconstructing relevant aspects of its user interface. In particular, we will see how: a console window collects all people, artifacts, and places in a manner akin to buddy lists; media items represent people and artifacts; places collect media items, where a person can adjust the salience of those items by various means; and places can be presented in several forms (e.g., as stand-alone windows, or within a places strip on the desktop border) to adjust a place's overall salience.

3.1 The Come Together Console

One of the most powerful features of Instant Messengers is the buddy list. In particular, a person can select their own small community of collaborators from a much broader population. The buddy list serves to provide awareness of others as online status, and as a way to selectively incorporated particular buddies into a conversation.

Come Together extends the affordances of a buddy list to include not only buddies, but artifacts and places. It does this via its main *console* window, which contains four tabs: People, Things, Places, and Me (Figure 1, a-d). Unlike Instant Messengers, each tab provides rich and alternative ways of viewing its contained items. The description below summarizes the basic features of these tabs. Due to space limitations, particulars of selecting and manipulating the items in tabs are not discussed, but include typical operations such as drag and drop, and selection of actions via context menus [16].

People tab. Figure 1a illustrates one person's view of the People tab, which is somewhat equivalent to a traditional buddy list. The bottom part (below the splitter) shows the community of all potential buddies as text, images and/or video (if enabled by the other person). Any person can be made into a buddy by selecting the 'I See' checkbox; when the other person reciprocates, the 'See Me' checkbox will be filled in (for example, Helen). The part above the splitter shows one's actual buddies. In this case, five buddies are visible across two rows: the lower row presents them as medium size image or video thumbnails, while the upper row shows them as small icons. The person can adjust the prominence of a particular buddy simply by dragging and dropping them between these rows. A buddy's on-line state is indicated by a grey scale (offline) or color (online) images or video.

Things tab. Figure 1b shows the Things tab, which collects all artifacts (e.g., web pages, documents, pictures, ongoing chats) viewable by that person. Three sections are visible. The top section includes one's own unshared things that are not part of

Figure 1. The Come Together Console Window, and its tabs

any place; the middle includes one's own things that are part of a place and thus sharable, while the bottom shows things created by others that are shared with that person because both are included in at least one place. Each section also has thumbnail and icon rows that let people adjust a thing's visual prominence. As with people, things can be represented by live thumbnails that show the current status of that thing, or as smaller abstract icons that represent the thing's type.

Places tab. Figure 1c shows one person's view of places. The top section shows the (three) titled places this person is part of. Each place shows a miniature of several people and things in it; the drop down at the right of each place will expand the view to reveal the other people and things in that place. The bottom section shows other places on the system that the person could potentially join; the idea is that giving other non-participants (i.e. those at the outermost periphery) a glimpse into these other places provides further opportunities for engagement. The button at the center lets a person create a new place; that new place would then contain that person. Other people can be invited into that place and existing or new artifacts can be added. Each place has a context menu that lets one of its occupants rename it, quit it, set it as private or public, and if public, set the fidelity of information visible to non-members.

Me tab. Figure 1d shows the Me tab, which lets a person control how others see them. The controls at its top let the person specify one's name and alter login credentials. Below that is a set of controls allowing the person to specify if he or she is viewable as a video feed (a preview is shown), and if so, what frame rate and camera to use. For cases when video is not being shown, the bottom controls let a person specify a static image (from file or from a video grab) to show instead; previews of that image at the icon and thumbnail sizes are displayed.

3.2 Come Together Media Items

Media items are containers of groupware content that allow people to drill down into that content to see its details [11]. Media items in Community Bar were well received by participants [15], as they proved an effective means for people to move from awareness to interaction (by drilling down), but proved less successful in terms of how people could adjust their salience by resizing them. Come Together re-implements media items in an altered form to keep its main strengths while overcoming its weaknesses.

From a software development point of view, a media item is a software plug-in: a container with a well-defined interface [12][16]. Thus CT is extensible: a developer can create a variety of different media items. CT is packaged with the following group-aware media items (new media item types could be created and added with routine software development effort [16]).

- *Presence item* shows a person's image or video;
- *Photo item* collects photos (e.g., by a person dragging / dropping them into the item) that can be viewed as a slide show;
- *Web item*, when given a URL, collects and displays that web page, where links can be followed as expected;
- *Chat item*, an example of a communication 'thing', serves as an instant messenger chat box, where it maintains a persistent transcript of the dialog;
- *Document item* displays a scrollable XPS document;
- *Screen item* (Figure 2) periodically captures and displays a live image of the screen of its owner to all viewers. The owner can adjust how often that image is refreshed, while other participants have different controls allowing them to adjust how they see the other person's screen.

Unlike Community Bar, a person using Come Together can create a media item both outside and within a place. If outside, it can be used as a stand-alone unshared item. That same media item can then (if desired) be brought into one or more places. These differences are critical. First, because a place is not required, media items can be used for individual work. Second, because a media item can be created within or incorporated into one or more places, the same item can become part of one or more group activities as opportunities arise.

CT media items extend the Community Bar's drill-down method of showing details of its content. CT uses the size of the item to gauge the amount of detail that would be appropriate, i.e., via 'semantic zooming'. A CT item's interface minimally demands that at least three key size thresholds are defined (the programmer is also free to have an item's semantic contents change continuously to match any size): the *icon view* is a size below 64^2 pixels, a *thumbnail view* is from 64^2 - 256^2, while a *full view* is any size larger than that. Each media item then decides what to display in that range. As well, CT media items are not limited to the sidebar approach of Community Bar.

Reconsider, for example, the people displayed in the People Tab in the two rows at the top of Figure 1a. Each person is actually displayed within a presence media item. The top row shows a 16^2 icon view of a person's static image. The next row displays presence items as 64^2 slow frame-rate live video (if available) or as higher resolution static images. For larger presence items, higher resolution full frame-rate video is displayed.

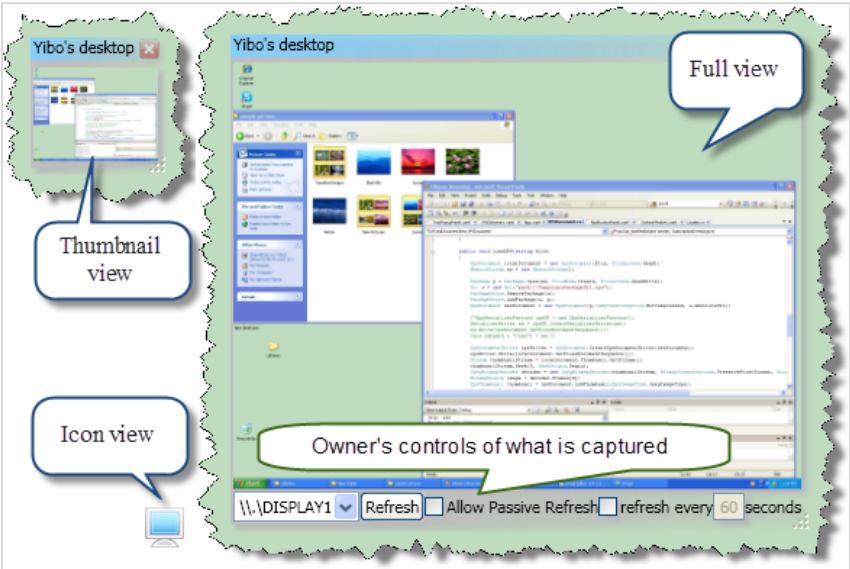

Figure 2. Icon, thumbnail, and full view of the screen item

Arifact-oriented media items work the same way. The screen item in Figure 2 shows: an abstract monitor icon in the icon view representing the media item type (bottom left); a shrunken live image of the screen in the thumbnail view (top left); and a scaled-to-fit screen image plus various controls in the full view (right). Both the web item and document item (e.g., as visible in Figure 1b, top) also have their semantic contents match their size: an abstract icon representing a web or document in the icon view, a shrunken image of the actual web or document page in the thumbnail view, and the actual web or document page in a web browser or document viewer in the full view. The chat item shows as much of its live content as can fit within its thumbnail view, and a scrollable transcript and input area in the full view (Figure 8b). The photo item shows an exemplar photo in its thumbnail view, and a controllable slide show in its full view (Figure 6c-e).

As we will see, Come Together uses media items' semantic zooming approach to offer details on demand, and to adjust an item's salience in a variety of ways. For example, dragging an item from the icon to thumbnail row changes its size and thus the richness of information displayed (e.g., as in Figure 1a-b). If a person hovers the mouse over the icon view, a thumbnail / text description will be displayed in a manner akin to a tooltip, e.g., Figure 3's document icon's tooltip. For full details, a media item instance can be dragged out of the console onto the desktop, where it appears in its own window. When a person resizes the media item window, its view will be adjusted as described above. Media items can also be part of a place and/or desktop border, where it will exhibit variations of its behavior.

3.3 Come Together Place Window

As previously discussed, if lightweight meetings are to work, creation of a group meeting place and bringing people and things into that place must be extremely easy. Thus a crucial component in Come Together is how a place is created. Come Together provides several methods to do this.

One way to create a place is by joining media items together via dragging and dropping. Figure 4 shows an example sequence, where Saul wants to show off his photos from his last camping trip. From the 'My Buddies' portion of his People tab in the Console (Figure 4, top left), Saul sees that his student Yibo is online. He drags outs Yibo's media item from the People tab onto the desktop (Figure 4, top). Saul then switches to the Things tab, and drags out his previously created photo item containing his camping photos close to Yibo's media item (Fig. 4 mid- left); the system recognizes this as a request to create a place, and draws a red bounding box around the two items indicating this (Figure 4, mid-right). Saul then drops the media item, and the place with the default name 'New Place' is automatically created containing Yibo and the photos (as thumbnails), and Saul (as an icon) (Figure 4, bottom left). As this is happening, Yibo gets a notification on his display telling him that he has been added to this place, where it appears as a popup 'toast' containing an icon view of the contents of that place (Figure 4, bottom right). If Yibo does nothing, Yibo would see this place on his Places Tab of his Console, i.e., he would be on the periphery of this place where he could go into it later. Alternately, he could click on that toast, and the place would open up on his desktop (similar to Saul's view at the bottom left of Figure 4), i.e., Yibo would move toward the center. What is important is that there is no formal invitation / acceptance mechanism: Saul's addition of Yibo into the place suffices to give Yibo full access to enter that place at any time. From this point on, new people or thing media items can be created within that place via: the place's context menu (not shown); by dragging and dropping other media items onto it; and even by dragging and dropping documents of certain types into the place. For example, dragging a picture into a place will automatically create a photo item around it, after which more photos can be dropped onto it. Similarly, dropping a web URL and an XPS document onto the place will create a web item or document item in that place. Media items can also be removed from a place.

There are other ways to create places. A second method allows a person to explicitly create a place from a stand-alone media item via a menu option. Figure 5 illustrates this, where Yibo raises the context menu atop of Helen's presence item. Through this, Yibo can create a place containing only himself and Helen, or can

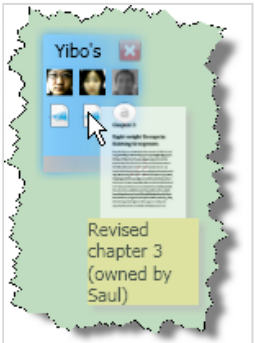

Figure 3. Thumbnail as an icon's tooltip

Figure 4. Creating a place by bringing together two media items

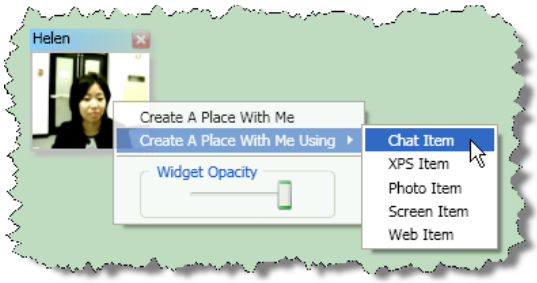

Figure 5. Creating a place from a media item

choose to also immediately populate that place with other (empty) media items via the context menu. For example, selecting the Chat item creates a place containing Helen, Yibo and a text chat, i.e., it is equivalent to an instant messenger conversation (with added features such as persistence and video, etc.). As before, Yibo (and now Helen) can drag existing media items (people or things) into that place at any time as the meeting unfolds. A third method varies the above, where a person can create a place by dragging a URL or a file of a particular type over a stand-alone presence item. For example, if Yibo dragged an XPS document file atop of Helen's presence item, CT would create a document media item containing that document, and would automatically create a place containing both Helen's presence item and the document item. As yet another method, a person can create a place directly within the console Places tab (using a button), after which that place can be dragged out onto the desktop and populated as desired.

An important difference from many groupware systems is that a Come Together place can be created that only contains a single person; there is no requirement that other people have to be invited into that place. Thus a place can not only serve for groupware meetings, but also as a way to define (and perhaps set aside for later perusal) a person's individual activities within a persistent context, e.g., as an activity as according to Activity Theory [14]. Other people can be brought into that place at a later time, but only if desired.

3.4 Exploring Media Items within a Place

Places include rapid interaction techniques to let a person quickly adjust the prominence of its contained media items, and / or to quickly drill down into details of a media item. The mechanism is a variation of the Tooltip Grande popups found in Community Bar [11] and SideShow [1].

When a place is created, by default it shows only the icon and thumbnail views of media items, arrange in two rows (e.g., as in Figure 6a). As mentioned, the person can drag and drop items between these rows to switch between icon and thumbnail representations, or raise a thumbnail tooltip from the icon (Figure 3), or even drag items around to customize the order of thumbnails or icons. However, a person can also see the full view of a media item quickly as a transient Tooltip Grande, i.e., within a large pane temporarily included in the place. Figure 6 shows examples. In Figure 6a, Saul is about to move his cursor over Yibo's presence item thumbnail. When it is over it, the place automatically increased in size to reveal a pane containing the full view of Yibo's media item (Figure 6b), that is, a larger video of Yibo running at a higher frame rate. In Figure 6c, Saul continued moving his cursor over the photo item thumbnail, and as a

a) Default view of a place

b) Presence tooltip grande

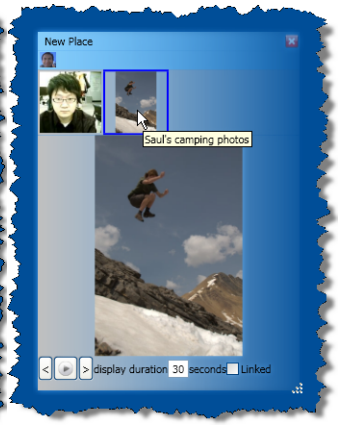

c) Photo tooltip grande

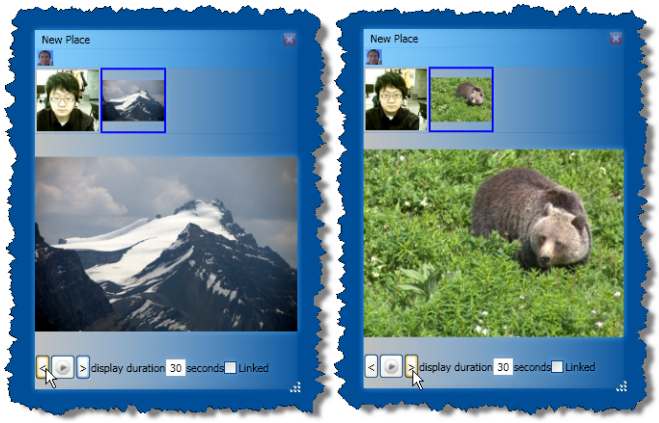

d,e) Locking down the photo tooltip grande to view photos

Figure 6. Exploring a place's media items

consequence the full view of the photo item replaced the video item. Saul now sees large photos cycled from that collection (Figure 6c). When he moves the cursor off the thumbnails, the full view disappears and the place shrinks down to its original small size (as in Figure 6a). In Figure 6d, Saul has locked down the larger views of the photo item by clicking on its thumbnail; the full view persists, which lets him move his cursor elsewhere. For example, he can now use the controls at the bottom (Figure 6d,e) to flip through the photo set, to adjust and play the photos as a slide show, or even to adjust how other's view of these photos are linked, i.e., as loosely coupled (where each can flip through the photos independently) or tightly coupled (where all see the currently selected photo). Saul can also drag and resize a media item out of the place and onto the desktop; while still part of the place, this allows one or more media items to be displayed in parallel at various sizes.

Collectively, this rapid drill-down and/or customization of the level of views of items within a place support a person's ability to adjust items within a place, i.e., to match one's personal center / periphery view of people and things within a place, and to easily move from awareness to interaction.

3.5 Places on the Desktop Border

A unique interface element of Come Together is the *desktop border*, which collects a place and its media items as a vertical or horizontal strip along the edges of the display. Its purpose is to let an end user further customize their individual view of their places. In particular, a person can move particular places even further to the periphery, while still allowing them to receive subtle notifications of ongoing changes to a place and opportunity to explore elements of a place. As we will see, the desktop border acts as a docking station for places, somewhat akin to the sidebar used in SideShow [1] and CommunityBar [11]. While the desktop border is initially empty, a person can drag and drop places onto that border at any time, and then adjust the size of media items within it to fine-tune an item's salience.

Figure 7 illustrates a segment of one person's desktop border containing two places, joined by a *connector*, titled "Yibo's new chapter" and "Wikipedia's page for groupware". Both places currently show that the person has a hybrid view onto its items. In the Wikipedia place, he wishes to monitor (albeit somewhat at the periphery) the online status of two of the people within it, as well as their ongoing conversation. Thus he has resized those three into a larger thumbnail view but kept the other two as icons. In this case, If he wishes greater involvement (i.e., to move towards the center), he can drag a place out of the desktop border onto the display, where it will be re-created as a place window.

More generally, the docking station supports two presentation levels—the icon view and the thumbnail view—of its media items. When shown as a horizontal strip (as in Figure 7), the upper part of the docking station stacks media items as icons, while the lower part arrays media items in their thumbnail views. The end user can drag a media item between these two parts to change its representation. However, the person can also move the cursor over a thumbnail and customize its size via the scroll wheel, e.g., to decrease the thumbnail size. As well, a person can double click the title bar to toggle between a compact icon-only view and the hybrid icon / custom thumbnail view. As with place windows, the person can drag out the media item onto the desktop and size it to whatever view they wish. Finally, if an item's content is changed significantly (as determined by the type of media item), it is outlined in red to indicate that its contents have changed since the person last hovered over it (this is also true in other views of media items within a place).

A person can also customize (albeit loosely) the relationships between places on the strip. For example, a place can be moved adjacent to another place to relate them as a cluster. One such cluster of two places is illustrated in Figure 7. This provides a (limited) sense of *civic structure* as suggested by the Locales Framework. However, this is an individual view; only the end user that created this adjacency pair will see that.

4. Example Scenario

We now illustrate how the above works via a deliberately simple example scenario. Yibo decides to use Come Together to capture, store and review some of his thesis source materials that he found on the web. He begins by capturing the Wikipedia 'Collaborative software' entry, as it contains a nice summary of groupware systems. He adds this web page to his Artifact collection using the Console, where he simply drags the URL of the page from the browser's address bar and drops it onto "My Private Things—Not

Figure 7. The Desktop Border

Shared In Any Place" section under the Things tab (Figure 8a). After Yibo confirms this, the web media item with the thumbnail of the web page appears (Figure 8a, right). Using the context menu, Yibo renames that media item to something more descriptive (e.g., "Wikipedia Groupware"). He then drags the web item onto the desktop (not shown), where he sizes it as a thumbnail to remind him to read that page. Later, he resizes it as a full view (e.g., similar to Figure 2), and he reviews that page by scrolling and navigate its links. What this sequence illustrates is that artifacts can be used by individuals and can exist outside of places; i.e., artifacts need not be shared with others before they can be created and used.

Yibo suddenly realizes that some of his colleagues may find this (and other groupware references he will collect) useful, so he decides to create a place that collects this and future references. He notices in the Console that Miaosen and Helen are both online, and includes them in this place by dragging out their presence items (showing their video feeds), and the web item together onto the desktop (Figure 8b, top left). A place appears containing these three thumbnail items and Yibo's icon (Figure 8b, top right). He renames "New Place" into "Wikipedia's page for groupware", adds a chat media item, and posts a message into it (Figure 8b, bottom left). Because Yibo is familiar with that web page, he changes its representation from a thumbnail to an icon (Figure 8c, top), but keeps the people and chat as thumbnails as he wishes to monitor them. Finally, he puts the whole place on the window border, as he is moving onto other things and wants this at the periphery of his attention. Helen and Miaosen can now enter that place, review that web page, add to the conversation, add other things to that place, and even invite new members (not shown). What this sequence illustrates is the highly fluid, customizable yet still lightweight nature of this group meeting place: how it can be created on the fly, how its contents appear, how others can see and use it.

5. Discussion

To summarize, Come Together's design embodied the following main concepts, which contribute over prior systems. All concepts cover the elements introduced in the introduction, as well as the primary notions suggested by the Locales Framework.

1. Locale formation via places is lightweight.
   - Places can be created implicitly simply by bringing two or more media items (or documents) together or via context menus attached to media items.
   - Places can be created on the fly. No prior configuration of the site and/or means is required. They need not be named, nor do their contents have to be specified *a priori*.
   - The means (artifacts and people) can exist outside of places and brought into it as needed.
2. Membership to places is lightweight.
   - Anyone can add a person to a place, where that person is located at the periphery and can move to the center as

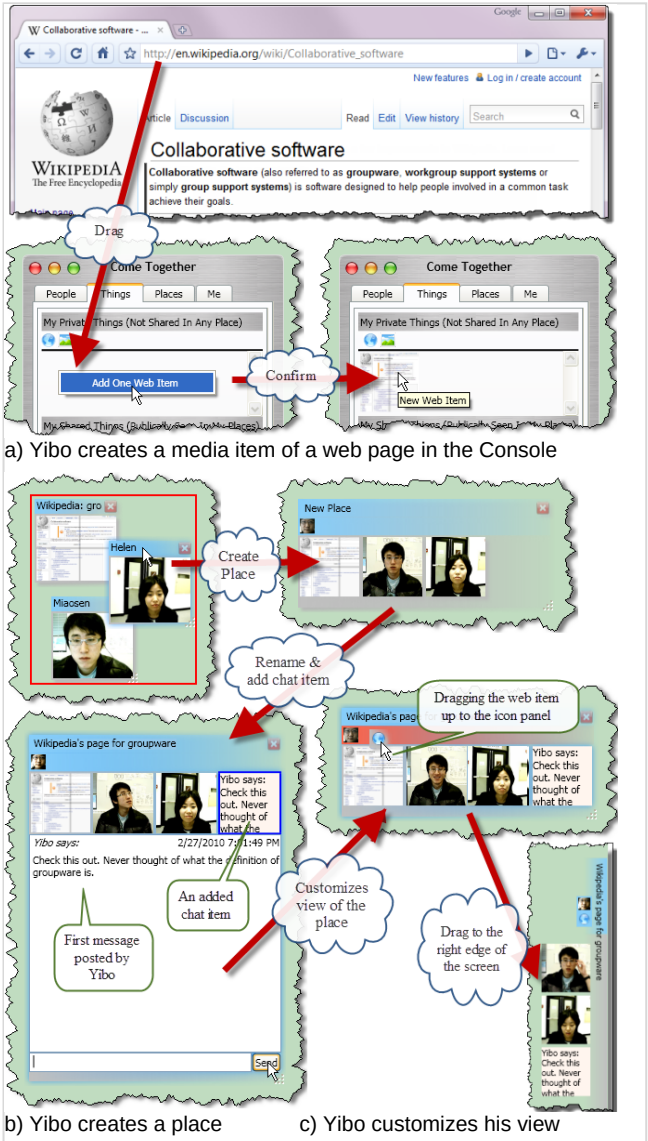

a) Yibo creates a media item of a web page in the Console

b) Yibo creates a place c) Yibo customizes his view

Figure 8. Sequences in a Scenario

desired. No formal invitation, acceptance, or moderation protocol is needed.

- A person can join an existing place simply by dropping their presence item atop of it.

3. Crafting the means within a place is lightweight.
   - Anyone can add new or existing artifacts to a place.
   - Artifacts that exist as either stand alone entities or within other places can also be brought into a place. That is, a single artifact can populate 0, 1 or more places, where changes made to that artifact in any place will propagate to other places.
   - Because artifacts and people media items are treated identically, people use the same interaction methods to add people and things. This reduces interface complexity.
4. Mutuality (or awareness) is lightweight.
   - The Console provides a 'one-stop' shopping view of the status of all key components (places, things, and people).
   - Places within the Console, or as place windows, or within window borders provide varying views of the activities within a place. Thumbnails show the live state of their items, and highlighting of items indicate what has changed since the person was last active within that place.
   - People can quickly drill down between icon/thumbnail and full view (via transient popups, pinning, or dragging out media items to the desktop) to move from awareness to interaction, and to see details on demand.
   - Notifications (as toasts) are used to indicate major changes within a place to others (only the addition of a member to a place was mentioned previously, but other notification events exist).
5. Moving oneself to reflect one's interest in the center/periphery spectrum is lightweight.
   - People can choose different views (and thus different levels of prominence) of a place, which in turn reflects how interested they are in that place, e.g., within the Console only, within the window border, and/or as a small or large place window.
6. Adjusting individual views of people and things on the center/periphery spectrum is lightweight.
   - The Console lets people rapidly select buddies from a broader pool of people, i.e., their individual view of who is important to them.
   - People and artifacts, which are treated identically in media items, can be rapidly switched between three presentation levels (icon, thumbnail and full view).
   - The particular type of place provides different ways to tune the size of items (e.g., thumbnail sizes on places in the windows border are adjustable, and media items dragged onto the desktop can be resized at will).
   - The order of items within a place is also customizable.
   - While not mentioned previously, a person can also adjust the transparency of stand-alone media items and places windows, i.e. to mute their appearance even further.
   - Each user's view does not affect how others may see a place or its media items.
7. Social practices, as listed in Table 1, are lightweight.
   - Come Together does not attempt to impose a social protocol onto the group. It assumes that people within the group can develop and follow their own social practices, e.g., their structure, culture, group formation and termination, membership, and roles.
   - Every participant in a place enjoys the same privileges. Every member can invite people into a place, add artifacts to a place, modify the place, and reconfigure its settings.
   - While there is some degree of access control to a place (e.g., what others can see via the console), it is up to the people within a place to decide who can be invited into it. Once invited, that person becomes an equal member.

6. Early Evaluation

We recruited 12 participants knowledgeable in HCI. Each walked through various short CT use cases, where she/he saw and experienced CT's features. Each individual was also given a review of CT's design goals and how CT's interface tried to meet that goal. Based on both, the individual was asked to provide his/her general impressions of system, as well as critique the specific features demonstrated in the walkthrough. Overall, we gathered quite positive reactions along with suggestions for improvements (see [16] for full details). Positive aspects included most of the way CT was designed around the Locales Framework,

for example, the light-weight formation of persistent locals, the easy way people could view a place's contents, the simplicity of adding a person to a place with no formal invitation process, the easy adjustment of awareness and involvement in a place. Areas for improvement included routine interface issues, e.g., a simpler way to leave a place, and the UI complexity of the console. More seriously, while people enjoyed CT's use of social protocols, some were concerned about what would happen in communities where people were not intimate collaborators. In those cases, they thought access control and invitation processes would be warranted. This is a tradeoff, as it would make place formation more heavy-weight.

The study affirms that our design goals for CT are all on the right track. However, the study also suggested aspects that require fine-tuning. We also have to be careful about the assumptions made of the expected community of users. As a consequence defaults should not be universally applied, but should be suggested by the site and/or end user. Ultimately, we need to experiment with alternative user interfaces and consider their design trade-offs. This is fairly typical of iterative design. Our initial design passed the overall test, where it appears to match our design goals. Fine-tuning will make it even more usable and flexible.

Cavaeat. The system has not been deployed for field testing, for a variety of reasons. It is still somewhat buggy (due to it pushing various features of Windows WPF and because of occasional race conditions present in its underlying distributed architecture). The groupware features within it are limited. For example, we have not implemented voice over IP, which we believe is crucial for real-time interaction. While out of band voice could be used (phone, Skype), setting those channels up would take away from the lightweight nature of the group activities. Other more functionally-rich media items need to be included. For example, we expect many users would want to share and edit their Microsoft Word or Powerpoint documents; our system currently handles only XPS documents. Thus deployment and study would be premature, as it would say more about how people react to minor interface issues, bugs, the limited tools available, and performance bottlenecks rather than its main design concepts. Our next steps are to make the system more robust, to do some redesign as mentioned above, then deploy and test it in actual use.

7. Status and Conclusions

Come Together is a fully working prototype available at http://grouplab.cpsc.ucalgary.ca/cookbook: it has all the features (and more!) described here. Full details concerning its implementation are in [16]. It runs on a Windows platform, and requires no special equipment or software except for a webcam.

Our approach in this paper was to argue for the various design features of a place-based system supporting lightweight group meetings, as informed by the Locales Framework and by considering the successes (and weakness) of two systems: Instant Messengers and Community Bar. Our contribution is largely in design, i.e., how we created and exploit notions such as Consoles (as a buddy list for people, things and places), media items, places, and the varying highly customizable views of all of them.

Acknowledgements. Funding was partially provided by the NSERC/iCORE/SMART Chair in Interactive Technologies. Steve Voida contributed insights in the project's early stages.

8. References

[1] Cadiz, JJ, Venolia, G.D., Jancke, G., and Gupta, A. (2002) Designing and deploying an information awareness interface. *Proc ACM CSCW*, 314-323.

[2] Dourish, P. (2006). Re-Space-ing Place: Place and Space Ten Years On. *Proc. ACM CSCW*, 299-308.

[3] Dourish, P. and Bellotti, V. (1992). Awareness and Coordination in Shared Workspaces. *Proc ACM CSCW*, 107-114.

[4] Edwards, K.W., Newman, M., Sedivy, J., Smith, T., Balfanz, D., Smetters, D., Chi Wong, H. and Shahram, I. (2002) Using Speakeasy for Ad Hoc Peer-to-Peer Collaboration. ACM CHI 2002.

[5] Fitzpatrick, G. (2003)*The Locales Framework: Understanding and Designing for Wicked Problems.* Kluwer.

[6] Greenberg, S. and Buxton, B. (2008) Usability Evaluation Considered Harmful (Some of the Time). *Proc. ACM CHI*, 111-120.

[7] Greenberg S. and Roseman, M. (2003) Using a Room Metaphor to Ease Transitions in Groupware. In M. Ackerman, V. Pipek, V. Wulf (Eds) *Sharing Expertise: Beyond Knowledge Management*, 203-256, MIT Press.

[8] Gutwin, C. and Greenberg, S. (2002) A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, 11(3-4):411-446

[9] Harrison, S. and Dourish, P. (1996). Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems. *Proc. ACM CSCW*, 67-76.

[10] Kraut, R., Egidio, C. and Galegher, J. (1990) Patterns of Contact and Communication in Scientific Research Collaboration. In *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work.* LEA Press, 149-181

[11] McEwan, G. and Greenberg, S. (2005) Supporting Social Worlds with the Community Bar. *Proc. ACM GROUP*, 21-30.

[12] McEwan, G., Greenberg, S., Rounding, M. and Boyle, M. (2006) Groupware Plug-ins: A Case Study of Extending Collaboration Functionality through Media Items. *Proc 2nd Int. Conf. Collaboration Technologies - CollabTech 2006*, IPSJ SIG Groupware and Network Services, 42-47, July.

[13] Nardi, B.A., Whittaker, S., and Bradner, E. (2000) Interaction and Outeraction: Instant Messaging in Action, *Proc. ACM CSCW*, 79 89

[14] Nardi, B. (1997). *Context and consciousness: Activity theory and human computer interaction*. MIT Press.

[15] Romero, N., McEwan, G. and Greenberg, S. (2007) A Field Study of Community Bar: (Mis)-matches between Theory and Practice. *Proc ACM Group*, 89-98.

[16] Sun, Y. *Groupware Places for Lightweight Meetings.* M.Sc. Thesis, Dept. Computer Science, Univ. Calgary, Canada.

[17] Tee, K., Greenberg, S. and Gutwin, C. (2009) Artifact Awareness through Screen Sharing for Distributed Groups. *Int. J. Human Computer Studies - IJHCS*, 67:677-702, Sept.

[18] Voida, S., Edwards, W.K., Grinter, R.E. & Ducheneaut, N. (2006). Share and share alike: Exploring the user interface affordances of file sharing. ACM CHI, 221–230.

[19] Whittaker, S., Frolich, D., and Daly-Jones, O. (1994) Informal workplace communication: What is it like and how might we support it? *Proc ACM CSCW*, 131-138.